

Route-Optimised Multicast Traffic for a Mobile Network Node

Description

Field of the invention

- 5 This invention relates to route-optimised multicast traffic for a mobile network node.

Background of the invention

- Traditional mobility support aims to provide continuous Internet connectivity to mobile hosts; thus offering host mobility support. In contrast,
10 network mobility support is concerned with situations where an entire network changes its point of attachment to the Internet topology and thus its reachability in the topology. Such a network in movement can be called a Mobile Network.

- There exist a large number of scenarios where such Mobile Networks
15 exist. Two out of many examples are:

- A Personal Area Network (PAN, i.e. a network of several personal devices attached to an individual) changing its point of attachment to the Internet topology while the user is walking in a shopping mall.
- A network embedded in a vehicle, such as a bus, a train or an aircraft
20 providing on-board Internet access to passengers. A passenger may use a single device (such as a laptop computer) or own a Mobile Network (such as a PAN), which then illustrates a case of a Mobile Network visiting a Mobile Network (that is to say nested mobility).

- As such, a Mobile Network can be defined as a set of nodes (so called
25 Mobile Network Nodes or MNNs) forming one or more IP-subnets attached to a Mobile Router (MR), the Mobile Network (the MR and all its attached MNNs) being mobile as a unit with respect to the rest of the Internet. Internet-Draft draft-ernst-monet-terminology-00.txt [Thierry Ernst, Hong-Yon Lach, "Network Mobility Support Terminology", draft-ernst-monet-terminology-00.txt, February

- 2 -

2002, work in progress] defines terminology for Mobile Networks that will be used in the following. Especially the following terms are defined:

- 5 • Local Fixed Node (LFN): A node permanently located within the Mobile Network and that does not change its point of attachment. A LFN can either be a LFH (Local Fixed Host) or a LFR (Local Fixed Router).
- 10 • Local Mobile Node (LMN): A mobile node that belongs to the Mobile Network and that changes its point of attachment from a link within the mobile network to another link within or outside the Mobile Network (the home link of the LMN is a link within Mobile Network). A LMN can either be a LMH (Local Mobile Host) or a LMR (Local Mobile Router).
- 15 • Visiting Mobile Node (VMN): A mobile node that does not belong to the Mobile Network and that changes its point of attachment from a link outside the Mobile Network to a link within the Mobile Network (the home link of the VMN is not a link within the Mobile Network). A VMN that attaches to a link within the Mobile Network obtains an address on that link. A VMN can either be a VMH (Visiting Mobile Host) or a VMR (Visiting Mobile Router).
- 20 • Mobile Network Prefix: A bit string that consists of some number of initial bits of an IP address which identifies a Mobile Network within the Internet topology. Nodes belonging to the Mobile Network (i.e. at least MR, LFNs and LMNs) share the same IPv6 "network identifier". For a single mobile IP-subnet, the Mobile Network Prefix is the "network identifier" of this subnet.
- 25 • Egress Interface of a MR: The interface attached to the home link if the Mobile Network is at home, or attached to a foreign link if the Mobile Network is in a foreign network.
- Ingress Interface of a MR: The interface attached to a link inside the Mobile Network.

30 Whereas the draft Mobile IPv6 specification [D. Johnson, C. Perkins, J. Arkko, "Mobility Support in IPv6", draft-ietf-mobileip-ipv6-20.txt, January 2003, work in progress] defines two means for a Mobile Node to receive multicast traffic while on the move, namely bi-directional tunnelling and

- 3 -

remote subscription, only the bi-directional tunnelling approach is currently foreseen in the case of a Mobile Network. In fact, most advanced proposals rely on bidirectional tunnelling between the Mobile Router and its Home Agent through which unicast and multicast traffic of Mobile Network Nodes should
5 be forwarded in both directions. Especially in the case of multicast traffic:

- Inbound multicast packets for MNN (i.e. multicast packets addressed to a multicast group G to which MNN has subscribed – MNN is a multicast receiver) are routed along the multicast tree in the backbone towards the Mobile Router's home link; intercepted by the MR's Home
10 Agent HA that tunnels them through a unicast tunnel to the MR, de-tunnelled by MR and forwarded along the multicast tree within the Mobile Network, and finally received by MNN as shown in Figure 1 of the accompanying drawings.
- Outbound packets (i.e. multicast packets sent by MNN to a multicast
15 group G – MNN is a multicast source) are routed towards the Mobile Router, reverse-tunnelled by MR to its Home Agent HA, and from there routed towards the multicast delivery tree as shown in Figure 2.

This mechanism does not provide route optimization to the MNNs since multicast packets between the multicast delivery tree (in the backbone) and
20 the MNN must go through the bi-directional tunnel between MR and HA, which potentially introduces a much longer path (take as illustrative example a MR deployed in a plane flying over the USA while its HA is located in France).

Thus there is a need for a means to enable MNNs to receive multicast traffic along an optimized path, that is to say, to have packets delivered
25 through the multicast tree to or from the current location of the Mobile Router without needing to transit through the MR Home Agent HA.

US Patent specification 20020150094 proposes a new IP multicast routing protocol, called "Hierarchical Level-based IP Multicasting" (HLIM) which is said to support not only host mobility (movement of IP hosts) but also
30 network mobility (movement of IP routers with or without attached hosts). Especially, HLIM is claimed to preserve on-the-shortest-path delivery of

- 4 -

multicast traffic from a source to a receiver located within a mobile network as the network changes its attachment point in the topology. However, HLIM, which has been designed for tactical networks, can only operate in very specific network topology (hierarchical networks), which is not the case of the Internet, thus limiting its applicability for commercial applications. In addition, HLIM requires all routers in the topology to run this new protocol which is unrealistic in the Internet whose multicast model is based on many multicast domains (owned by different parties) and possibly running different multicast protocols (such as DVMRP, MOSPF, PIM-SM, PIM-DM, CBT, for example).
5 Thus HLIM does not provide a means to support route optimised delivery of multicast traffic to a mobile network roaming in the Internet, irrespective of the multicast routing protocols used within and outside of the mobile network.
10

It is not desirable for a Mobile Router to rely on relaying multicast routing signalling messages (used to manage the multicast tree) between the nodes in the mobile network and the visited network (instead of through its home network and its Home Agent HA) in order to reconstruct a branch of the multicast tree towards the current location of the multicast-enabled mobile router. This approach is applicable if and only if the same multicast routing protocol is run both within the mobile network and visited network. As explained above, due to the very large number of existing multicast protocols, this requirement will rarely be met in practice. As a result, this approach does not enable route-optimised delivery of multicast traffic irrespective of the location of the mobile network in the Internet. In addition, in practice, security policies of the visited network will generally forbid any injection of routing signalling (unicast and multicast) from non-authorized nodes such as a visiting mobile router (the mobile router may be owned by a different organisation).
15
20
25

It has been proposed that the Mobile Network deploy on all routers within the mobile network a mechanism called "IGMP/MLD-based Multicast Forwarding" [B. Fenner, H. He, Nortel Networks, B. Haberman, H. Sandick, "IGMP/MLD-based Multicast Forwarding ("IGMP/MLD Proxying"), draft-ietf-magma-igmp-proxy-02.txt, March 2003, work in progress] instead of running a multicast routing protocol internally. This approach is intended to allow the
30

- 5 -

Mobile Router to collect all multicast group membership information coming from within its mobile network, and subscribe itself to all those multicast groups using IGMP/MLD protocol with the multicast-enabled access router in the visited domain. Group membership information will be relayed hop-by-hop, in the mobile network, from the intended multicast receiver up to the Mobile Router, by means of all intermediate fixed routers proxying incoming IGMP/MLD Report messages received towards its parent router (this is known as IGMP proxying, or MLD proxying). In this approach, the Mobile Router handles the multicast subscription in the visited domain on behalf of all the nodes in the mobile network. Upon movement, it will trigger reconstruction of a new multicast branch at its new location by sending MLD Reports to its new attachment point. However this approach requires heavy manual configuration, in particular to define upstream and downstream interfaces, on each router in the mobile network to make its internal topology like a tree routed at the Mobile Router. This makes this approach applicable only for relatively small mobile networks with stable internal topology. In addition, it imposes deployment of a new forwarding mechanism (IGMP/MLD proxy) on each internal router, and does not support route-optimised delivery of multicast traffic for any other form of multicast routing deployed in the Mobile Network. This is a limitation, especially for large mobile networks where regular multicast routing protocols are expected to be deployed to ease multicast support within the mobile network.

Thus there is a need for a mechanism enabling route-optimised delivery of multicast traffic to and from a mobile network:

- irrespective of the location of the mobile network in the Internet,
- irrespective of the type of multicast routing protocols used within and outside of the mobile network, and
- through extension of the Mobile Router involved alone, that is to say with no change to any node in the mobile network nor in the Internet.

- 6 -

Summary of the invention

The present invention provides a method of communicating traffic from a source to a group of nodes including a Mobile Network Node (MNN) in a Mobile Network using one or more multicast protocols and apparatus for use
5 in such a method as described in the accompanying claims.

Brief description of the drawings

Figure 1 is a schematic diagram of routing of inbound multicast packets in accordance with a known method,

Figure 2 is a schematic diagram of routing of outbound multicast packets
10 in accordance with the method of Figure 1,

Figure 3 is a schematic diagram of routing of inbound multicast packets in accordance with one embodiment of the invention, given by way of example,

Figure 4 is a flow chart of steps in the method shown in Figure 3, and

Figure 5 is an example of a group list maintained in the method shown in
15 Figure 3.

Detailed description of the preferred embodiments

The embodiments of the present invention shown in Figures 3 to 5 of the accompanying drawings provide a large degree of route optimisation for
20 multicast traffic offering certain key advantages:

- Minimal delay since packets are sent on the shortest path. This is important since many multicast applications have stringent requirements in term of delay (for example audio/video streaming, audio/video conferencing).
- 25 • Reduced probability of packet loss due to congestion since packets are sent along a shorter path. Again for real time multicast applications, this will improve quality of the stream at the receiver side.
- Enhanced scalability and robustness. By bypassing the Home Agent HA of the Mobile Router, which may easily be overloaded due to the

- 7 -

concentration of multicast traffic at this point, route optimisation reduces the probability of failure.

- Reduced bandwidth overhead as packets are not tunnelled. This helps optimising network resources.
- 5 • Maximal PMTU (Path Maximum Transmission Unit) on MNN-CN path, which ensures minimal fragmentation of payload.

The embodiment of the present invention shown in Figure 3 comprises a Multicast Signalling Gateway (MSG) co-located with the Mobile Router and having an MSG-enabled network interface to achieve route-optimised delivery
10 of multicast traffic to the mobile network irrespective of the location of the mobile network in the Internet, and irrespective of the multicast routing protocols used within and outside of the mobile network.

A key principle of the MSG is to translate messages of a multicast routing protocol (MRP) into messages of a group membership protocol
15 (GMP). It will be appreciated that this functionality of the MSG is completely different from the known ways that MRP and GMP protocols interoperate, including the translation of GMP messages into MRP messages.

A possible way for the MSG to generate the GMP messages, as detailed below, is to rely on the so-called "service interface" provided by these GMP
20 protocols. The "service interface" can be used by upper-layer protocols or application programs to ask the IP layer to enable and disable reception of packets sent to specific IP multicast addresses (optionally only from a given set of sources). This service interface can be understood as a function call, typically that is made available at the socket API level. This is available for
25 both IPv4 (IGMP) and IPv6 (MLD).

Multicast Routing Protocols (MRP) are protocols responsible for the construction of a multicast delivery tree, for instance DVMRP, MOSPF, PIM-SM, PIM-DM, CBT, etc. Basically, two families of multicast routing protocols can be distinguished:

- 30 • Protocols using explicit signalling to build the multicast tree: Those protocols define specific messages to be used by a receiver's multicast

- 8 -

router to trigger establishment of a multicast delivery branch towards itself. PIM-SM is one example of such protocols. These messages can be divided into two main categories:

- 5 ▪ "Join group": Messages used by a multicast router to join the multicast delivery tree. An example of such message is PIM-SM Join.
- "Leave group": Messages used by a multicast router to leave the multicast delivery tree. An example of such a message is PIM-SM Prune.
- 10 • Protocol using flooding: These protocols do not necessarily define specific messages to be used by a receiver's multicast router to join a multicast delivery tree. On the contrary, packets sent by a multicast source are flooded over the whole network and listened by any interested nodes. Most of existing protocols of this category however
- 15 include explicit signalling enabling to "prune" a branch to stop delivery of multicast traffic over a region of the network where there are no interested receivers. They also include support for "grafting" a branch when a new receiver appears. PIM-DM is one example of such protocols.
- 20 Group membership protocols (GMP) are protocols enabling a multicast receiver to announce its interest in receiving multicast packets sent to a multicast group G. Those are the Internet Group Management Protocol (IGMP) for IPv4, and Multicast Listener Discovery protocol (MLD) for IPv6. Note that recent versions of these protocols, IGMPv3 [B. Cain, S. Deering, B.
- 25 Fenner, I. Kouvelas, A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC3376, May 2002] and MLDv2 [R. Vida, L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", draft-vida-mld-v2-06.txt, November 2002, work in progress], when compared to the previous versions, add support for "source filtering". This refers to the ability for a receiver to
- 30 report interest in listening to packets only from specific source addresses, or from all but specific source addresses, sent to a multicast address.

- 9 -

By 'MSG-enabled network interface' is meant a network interface of a node running a multicast routing protocol on which Multicast Signalling Gateway operations are activated. A node may have several MSG-enabled network interfaces at the same time. In the case of the Mobile Router, all its egress interfaces should be MSG-enabled to achieve interworking between multicast routing protocols within and outside the mobile network thanks to a Group Membership Protocol.

Figure 3 illustrates the use of MSG for the case of a mobile router (MR) equipped with a single egress interface, by way of example. Multicast routing protocol MRP#1 is run within the mobile network. The MR is attached to a visited network that runs a different multicast routing protocol MRP#2. Both MRP#1 and MRP#2 are assumed to be part of the "explicit signalling" family. IPv6 is also assumed both in the mobile network and the visited network. Because of the IPv6 context, the group membership protocol (GMP) is MLD. MSG is enabled on MR's egress interface.

When a node MNN within the mobile network subscribes to multicast group G, it sends an MLD_Report(G) message of Group Membership Protocol towards his local multicast router LFR1. Since LFR1 is not yet attached to the multicast tree of group G (assuming MNN is the first receiver for group G below LFR1), LFR1 sends an explicit MRP#1_Join(G) message of Multicast Routing Protocol #1 within the mobile network to trigger establishment of a delivery branch towards LFR1. This branch establishment request propagates within the mobile network, eventually up to the Mobile Router whose local instance of MRP#1 protocol would decide to issue a MRP#1_Join(G) towards the egress interface. Since this interface is MSG-enabled, MR instead issues an MLD_Report(G) message of Group Membership Protocol towards Access Router (AR) in the visited network, which results in MRP#2_Join(G) messages of Multicast Routing Protocol #2 propagating within the visited network towards the multicast delivery tree of group G. These operations have enabled creation of two multicast delivery branches (one in each multicast domain) interconnected by the Mobile Router. As a result, multicast packets (e.g. from a source S) sent to group G

- 10 -

are MRP#2-native multicast routed towards MR in the visited network, and there from MRP#1-native multicast routed towards MNN.

The MSG-enabled Mobile Router MR handles the multicast subscription in the visited domain on behalf of all the nodes in the mobile network. The
5 MSG automatically discovers the subscription information (that is to say group and list of sources of interest) from the MRP messages arriving from within the mobile network. When MR changes its point of attachment to the visited network (or Internet), the MR will trigger reconstruction of one or more new
10 multicast branches at its new location by sending MLD Report(s) to its new attachment point for the multicast group(s) it has subscribed to. This is referred to as Remote Subscription.

Figure 4 is a flow chart of operations undertaken by the Multicast Signalling Gateway (MSG) in one embodiment of the present invention when detecting a Multicast Routing Protocol message (MRP message) about to be
15 sent through a network interface ifc_i of the node hosting the MSG.

If the interface ifc_i is not MSG-enabled, then the MSG just ignores the message.

If the interface ifc_i is MSG-enabled, then the MRP message is analysed to determine:

- 20 • The class of the MRP message: The class takes one of the following three values {JOIN, LEAVE, NONE} as a function of the type of the MRP message (or the impact it has on the multicast delivery tree). Any message of a MRP can be assigned one of these three values. This information can be stored within a table that is referred to as a class
25 table to be used by the MSG for the purpose of classifying messages of a given multicast routing protocol.
- The target of the MRP message: the target consists of the multicast Group G the message refers to, possibly together with an address of a source S. If a source address is not present, this means the packet
30 refers to any potential source. In this case, the wildcard "*" symbol is used to represent all sources: S=*. The target can be expressed as a

- 11 -

couple (S,G), or (*,G) in case no source specific information is conveyed.

If the class of the MRP message is NONE, no specific action is required from the MSG for this packet. This typically means that the packet has no semantic that should be translated into a GMP protocol message for the purpose of realising interworking between multicast protocols on both side of the MSG.

If the class of the MRP message is JOIN, the source S (from the target) must be added to the existing list of sources of group G (from the target) maintained by MSG for interface ifc_i: MSG_srclist(ifc_i,G). This is the list of sources of group G for which the MSG should maintain reception of traffic, through interface ifc_i. For this purpose the MSG maintains a list that is referred to as a group list, as illustrated in Figure 5, that includes, for each MSG-enabled interface, the list of groups that are of interest together with their respective lists of sources. Once MSG_srclist(ifc_i,G) has been updated (or created in case of a new entry) with the source S from the target, it should be checked whether this adding of S has modified MSG_srclist(ifc_i,G). If MSG_srclist(ifc_i,G) has not been modified then no action is required. On the other hand, if MSG_srclist(ifc_i,G) has changed then the MSG must renew the GMP subscription for this new set of sources of group G on interface ifc_i. The MSG can use the GMP "services interface" (or API) for this purpose as any other multicast application does.

If the class of the MRP message is LEAVE, the source S (from the target) must be removed from the existing list of sources of group G (from the target) maintained by MSG for interface ifc_i: MSG_srclist(ifc_i,G). Once MSG_srclist(ifc_i,G) has been updated, it should be checked whether this removal of S has modified MSG_srclist(ifc_i,G). If MSG_srclist(ifc_i,G) has not been modified then no action is required. On the other hand, if MSG_srclist(ifc_i,G) has changed and is now empty, the MSG must terminate the GMP subscription to group G on interface ifc_i. In addition, the MSG may remove the entry for group G in its group list for ifc_i. If the updated

- 12 -

MSG_srclist(ifc_i,G) is not empty, the MSG must renew the GMP subscription for this new set of sources of group G on interface ifc_i. The MSG can use the GMP "services interface" (or API) for this purpose as any other multicast application does.

5 The following arithmetic can be used when adding (+) or removing (-) a source (S or *) to/from a source list MSG_srclist(ifc_i,G):

- A source can appear only once in the list: $S + S = S$, $S - S = \emptyset$
- Adding all sources (*) to a list makes the list include all sources:
 $\text{srclist} + * = *$ (of course $* + * = *$)

10 • Removing all sources (*) from a list makes the list be empty (\emptyset):

$$\text{srclist} - * = \emptyset \text{ (especially, } * - * = \emptyset \text{)}$$

- Removing a defined source S from a list that includes all sources (*) does not change the list:

$$* - S = *$$

15 Figure 5 shows an example of group list maintained by a MSG. The following table shows the class table that can be used by a MSG for the PIM-SM multicast routing protocol. Similar class tables can be established for any multicast routing protocol (having explicit signalling) to be used by the MSG.

- 13 -

Class table for PIM-SM messages	
PIM-SM Messages	Class
Hello	NONE
Bootstrap	NONE
Candidate-RP-advertisement	NONE
Register	NONE
RegisterStop	NONE
(*,*,RP) Join	NONE
(*,*,RP) Prune	NONE
(*,G) Join	JOIN
(*,G) Prune	LEAVE
(S,G) Join	JOIN
(S,G) Prune	LEAVE
(S,G,rpt) Join	NONE
(S,G,rpt) Prune	NONE
(*,G) Assert	NONE
(S,G) Assert	NONE

Forwarding of multicast packets on an MSG-enabled node is very simple, and actually transparent to the MSG. Forwarding of multicast packets is done according to the multicast forwarding table of the MRP protocol run by the MSG-enabled node. This is true irrespective of whether the incoming interface is MSG-enabled or not.

Especially, in the case of the Mobile Network, multicast packets from a source external to the Mobile Network (that is to say, which have been subscribed to through the MSG-enabled interface) will be multicast-routed towards MR's egress interface (MSG-enabled) and therefrom multicast-routed within the Mobile Networks according to the local multicast forwarding table of the MR.

When the MSG-enabled interface of the Mobile Router (MR) changes its point of attachment to the visited network (or Internet), the MSG will trigger reconstruction of needed multicast branches at its new location by sending GMP report messages to subscribe to the multicast group(s) (and respective sources) listed in the group list for the MSG-enabled interface.

A Mobile Router (MR) for which a given egress interface is at home can decide either to configure this interface as MSG-enabled or not to do so.

- 14 -

Irrespective of the option selected, multicast traffic will be routed to MNNs in the same manner (according to the local multicast forwarding table of the MR). However, configuring the interface as MSG-enabled even when at home, will allow ongoing multicast communications to be maintained when the interface attaches to another topological location. This is because MSG will have learnt the list of ongoing groups G (and associated sources) and will then be able to re-subscribe to them at the new location.

In case MSG-enabled interface of the Mobile Router (MR) returns home and the Mobile Router (MR) decides to deactivate MSG on this interface, the multicast routing protocol will then operate normally through this interface towards the home network. In such a case, the MSG may remove any state in its group list for that given interface.

Several approaches can be taken for the implementation of the Multicast Signalling Gateway (MSG). In particular, for example, it may be implemented 1) as a separate software module or 2) as an extension (patch) of a multicast routing protocol (MRP) implementation.

In both cases, in this embodiment of the present invention, the MSG interface towards the Group Membership Protocol (MSG-GMP interface) can be implemented by relying on the existing GMP "service interface": MulticastListen(socket, interface, multicast address, filter mode, source list).

In this approach the MSG asks the IP layer (GMP) to enable and disable reception of packets sent to a specific IP multicast address in the same way as any multicast application program does (for example through a GMP-enabled socket API). Implementation based on the above service interface can be realised in two different ways:

- a) The MSG software handles the aggregation of sources for a given multicast group (MSG_srclist(ifc_i,G)) and uses a single socket identifier (sid) to pass the whole list:
 MulticastListen(sid, ifc_i, G, INCLUDE, MSG_srclist(ifc_i,G)), or
 MulticastListen(sid, ifc_i, G, EXCLUDE, {}), iff MSG_srclist(ifc_i,G)

== *

- 15 -

b) The MSG software uses several socket ids (one per (S,G) target) to pass the respective target (S,G) and associated class derived from the MRP message to the MLD module who is then in charge of aggregating the sources for a given multicast group. For a given target (S,G):

- 5 o If class is JOIN:
 - MulticastListen(target_sid, ifc_i, G, INCLUDE, S), or
 - MulticastListen(target_sid, ifc_i, G, EXCLUDE, {}), if S == *
- 10 o If class is LEAVE:
 - MulticastListen(target_sid, ifc_i, G, EXCLUDE, S), or
 - MulticastListen(target_sid, ifc_i, G, INCLUDE, {}), if S == *

This second approach requires the MSG software to generate a unique target_sid per target and store it into the MSG group list.

The implementation of the MSG-MRP interface towards the multicast routing protocol (MRP) will depend on whether approach 1) or 2) is chosen. The objective of this interface is for MSG to detect "MRP Message ready for sending on interface ifc_i" (see MSG state machine in Figure 4) to trigger MSG operations.

- Approach 1): MSG, as an independent software module, can detect MRP messages by monitoring packets sent over the interface ifc_i.
- Approach 2): MSG operations can be directly triggered by the MRP protocol implementation. For instance, the MSG-patched MRP implementation would invoke MSG procedures with the correct class and target information instead of sending a MRP messages through a MSG-enabled interface ifc_i.

In approach 1) MSG software is completely independent from the MRP and as a consequence may be used for any MRP as long as the corresponding class table is known. This eases software reusability.

- 16 -

In approach 2) MSG software is incorporated within the MRP implementation of the MSG-enabled node. This may provide more efficient implementation, for instance for detecting when MSG operation should be triggered.

5 While the Multicast Signalling Gateway (MSG) enables route-optimized delivery of multicast packets to a mobile network, it is also a very valuable approach to solve other types of issues. Below are two examples of other possible applications of the MSG:

- 10 • MSG for fast multicast domains interconnection: The MSG may be deployed as a temporary interworking point between two fixed leaf multicast domains, for instance in case of failure of the interconnecting multicast backbone. As such MSG offers a temporary failure-recovery solution that is fast and easy to deploy for network administrators.
- 15 • MSG for IPv4/IPv6 transition: The MSG offers a very convenient approach to interconnect IPv4 and IPv6 multicast clouds, so that IPv6 receivers can receive multicast traffic from IPv4 sources and vice versa. This is a very important use case considering that IPv4 and IPv6 protocols are foreseen to coexist for a relatively long time,
20 until the transition phase from all-IPv4 to all-IPv6 completes. For this purpose the MSG should translate multicast packets (together with unicast source addresses and multicast destination addresses of multicast packets) from IPv4 to IPv6 and vice versa. Most previously known mechanisms for IPv4/IPv6 transition apply only
25 for unicast traffic. A proposal has been made of a mechanism for multicast traffic [K. Tsuchiya et al, "An IPv6/IPv4 Multicast Translator based on IGMP/MLD Proxying (mtp)", draft-tsichiya-mtp-01.txt, February 2003, work in progress] that features an IPv4-IPv6 protocol translator and address mapper in order to realise
30 forwarding of multicast packets between the IPv4 and IPv6 domains (and vice versa). However, this proposal mandates that the

- 17 -

translator be manually preconfigured (by network administrator) to forward traffic for a given set of IPv4 (respectively IPv6) multicast addresses. During this setup, the translator joins ALL those preconfigured groups so that it is able to receive related traffic. This mechanism is clearly not efficient since it requires human intervention each time a new multicast group has to be translated. The MSG (extended with an IPv4-IPv6 header translator and address mapper) solves this issue by dynamically triggering reception of IPv4 (respectively IPv6) multicast traffic at the MSG only when needed by the IPv6 (respectively IPv4) domain.

Note that such use of an MSG-enabled Mobile Router enables an IPv6 (respectively IPv4) Mobile Network to roam into an IPv4 visited network (respectively IPv6) and receive multicast traffic from this visited network.

It will be appreciated that the embodiments of the invention described above offer a number of advantageous features. For example:

- a means to enable delivery of multicast traffic to MNNs along an optimal path irrespective of the location of the mobile network in the Internet.
- route optimisation for multicast traffic transparently to the MNNs: no change is required at the MNNs, even basic, non mobility-aware LFNs can benefit from route optimization (e.g. basic electronic devices in a car). Route optimisation is achieved by the sole MR.
- route optimised delivery of multicast traffic for nested mobile networks (i.e. a mobile networks visiting another mobile network), irrespective of the number of level in the aggregate of nested mobile networks
- mechanisms for seamless mobility of mobile multicast host based on the remote subscription approach applicable to the case of mobile networks.

- 18 -

The above features are applicable without any changes to existing multicast routing protocols when co-located on a Mobile Router to enable route-optimized delivery of multicast packets.

Additionally, the MSG is:

- 5
 - applicable to other scenarios such as IPv4/IPv6 transitions.
 - very simple to implement.
 - does not require any change to existing protocols nor extensions on any other nodes in the network.